


COPYRIGHT NOTICE:

Robin Wilson: Four Colors Suffice

is published by Princeton University Press and copyrighted, © 2003, by Princeton University Press. All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher, except for reading and browsing via the World Wide Web. Users are not permitted to mount this file on any network servers.

For COURSE PACK and other PERMISSIONS, refer to entry on previous page. For more information, send e-mail to permissions@pupress.princeton.edu



The four-colour problem

Before we embark upon our historical journey, there are a number of basic questions to be answered. First and foremost, of course, is this:

What is the four-colour problem?

The four-colour problem is very simply stated, and has to do with the colouring of maps. Naturally, when colouring a map we wish to colour neighbouring countries differently so that we can tell them apart. So how many colours do we need to colour the entire map?

At first sight it seems likely that the more complicated the map, the more colours will be required – but surprisingly this is not so. It seems that at most four colours are needed to colour any map – for example, in the map of Britain shown opposite, the counties have been coloured with just four colours. This, then, is the four-colour problem:

Four-colour problem

Can every map be coloured with at most four colours in such a way that neighbouring countries are coloured differently?

Why is it interesting?

Solving any type of puzzle, such as a jigsaw or crossword puzzle, can be enjoyed purely for relaxation and recreation, and certainly the four-colour problem has provided many hours of enjoyment – and frustration – for many people. At another level, the four-colour problem can be regarded as a challenge. Just as climbing a mountain can present a climber with major physical obstacles to overcome, so this problem – so simple to state, yet apparently so hard to solve – presents mathematicians with an intellectual challenge of enormous complexity.

Is it important?

Rather surprisingly, perhaps, the four-colour problem has been of little importance for mapmakers and cartographers. In an article on the problem's origin written in 1965, the mathematical historian Kenneth May observed:

A sample of atlases in the large collection of the Library of Congress indicates no tendency to minimize the number of colors used. Maps utilizing only four colours are rare, and those that do usually require only three. Books on cartography and the history of map-making do not mention the four-color property, though they often discuss various other problems relating to the coloring of maps . . .

The four-color conjecture cannot claim either origin or application in cartography.

Equally, there seems to have been no interest shown by quilt-makers, patchworkers, mosaicists and others in restricting the number of colours of patches or tiles to four.

However, the four-colour problem is more than just a curiosity. In spite of its recreational nature, the various attempts to solve it over the years have stimulated the development of much exciting mathematics with many practical applications to important real-world problems. Many practical network problems – on road and rail networks, for example, or communication networks – derive ultimately from map-colouring problems. Indeed, according to a recent book in the related area of graph theory (the study of connections between objects), the entire development of the subject can be traced back to attempts to solve the four-colour problem. In the theory of computing, recent investigations into algorithms (step-by-step procedures for solving problems) are similarly linked to colouring problems. The four-colour problem itself may not be part of the mathematical mainstream, but the advances it has inspired are playing an increasingly important role in the evolution of mathematics.

What is meant by ‘solving’ it?

To ‘prove the four-colour theorem’ is to show that four colours are sufficient to colour all maps – be they geographical maps of the world, or fictitious ones we may care to invent. If its statement were false, we would have to demonstrate the fact by presenting a map that requires five colours or more – just a single map will do. But if the statement is true, then we must prove it for *all possible* maps: it is not enough to colour millions or even billions of maps, because there may still be a map that we have missed whose

arrangement of countries requires five or more colours. In other areas of science, proofs assert that a given hypothesis is overwhelmingly likely given the underlying assumptions and the experiments that have been carried out, but a mathematical proof has to be *absolute* – no exceptions are allowed. To prove the four-colour theorem we must find a general argument that applies to all maps, and discovering such an argument turns out to require the development of a great deal of theoretical machinery.

Who posed it, and how was it solved?


As we shall see, the four-colour problem was first posed by Francis Guthrie around a hundred and fifty years ago, but more than a century of colouring maps and developing the necessary theoretical machinery would pass before it was established with certainty that four colours suffice for all maps. Even then, difficult philosophical questions have remained. The eventual solution, by Wolfgang Haken and Kenneth Appel in 1976, required over a thousand hours of computer time, and was greeted with enthusiasm but also with dismay. In particular, mathematicians continue to argue about whether a problem can be considered solved if its solution cannot be checked directly by hand.

PAINTING BY NUMBERS


Before we begin our historical narrative, we need to explain more clearly what the four-colour problem is, and what our underlying assumptions are. For example, what do we mean by a ‘map’?

We may think of a *map* as consisting of a number of *countries* or *regions*. These may be *counties* if part of a British map, or *states* if part of a map of the USA. The boundary of each country is made


up of several *boundary lines*, and these boundary lines intersect at various *meeting points*. Two countries with a boundary line in common are called *neighbouring countries* – for example, in the following map the countries *A* and *B* are neighbouring countries.




When colouring a map we must always give neighbouring countries different colours. Notice that some maps do need four colours: map (a) below has four mutually neighbouring countries, each meeting the other three; these countries must all be coloured differently, so four colours are required. This can happen in practice: in map (b), *Belgium*, *France*, *Germany* and *Luxembourg* are all neighbouring countries, so the map requires four colours.



Some maps do not need four colours. For example, in each of the following maps, the outer ring of countries can be coloured with two colours (*red* and *green*) that alternate as we proceed around the ring; the country in the centre must be then assigned another colour (*blue*), so that just three colours are required.



Notice that four colours may be needed even when the map does not contain four mutually neighbouring countries. In the map below, the outer ring of five countries cannot be coloured with two alternating colours, and so requires a third colour. The country in the centre must then be coloured differently from these three, so four colours are required.



This last situation also arises when we colour the 48 contiguous states of the USA (excluding Alaska and Hawaii). Here, the western state of *Nevada* is surrounded by a ring of five states, *Oregon*, *Idaho*, *Utah*, *Arizona* and *California*: this ring of states requires three colours, and *Nevada* must then be assigned a fourth colour. This colouring with four colours can then be extended to the entire map, as shown opposite.

We can make a couple of further observations about this map. Notice first that at one point of the United States four states meet – *Utah*, *Colorado*, *New Mexico* and *Arizona*. We shall adopt the convention that when two countries meet at a single point, we are allowed to colour them the same – so *Utah* and *New Mexico* may be coloured the same, as may *Colorado* and *Arizona*. This convention is necessary, since otherwise we could construct ‘pie maps’




that require as many colours as we choose – for example, the eight-slice pie map below would need eight colours, since all eight slices meet at the centre. With our convention, this map requires only two colours.



Another familiar ‘map’ that needs only two colours is the chess-board. At each meeting point of four squares we alternate the colours *black* and *white*, producing the usual chessboard colouring (see over).

The other feature of the United States map we need to note is that the state of *Michigan* is in two parts (separated by one of the Great Lakes) that must be assigned the same colour. Although a




divided country or state may cause no difficulty in particular situations, such as in the map of the United States, consider the following map in which the two regions numbered 1 are to be regarded as parts of the same country. Here, each of the five 'countries' has a boundary line in common with the other four, so five colours are required:




From now on, we avoid such undesirable situations by requiring that each country must be in one piece.


Some people also like to include the 'exterior region' in their colourings. Generally speaking, it is immaterial whether we do so or not, since we can regard this exterior region as an extra (ring-shaped) country (see opposite, top).



However, a more useful way of thinking about the exterior region is to consider the map as drawn on a globe, rather than on the plane. In this case, the exterior region appears as no different from any other region.



In fact, drawing maps on the plane is equivalent to drawing maps on a globe or, in the language of mathematics, on the surface of a sphere. We can see this from the following illustration, which depicts what is termed a *stereographic projection*, from the



sphere onto the plane on which it rests. Starting with any map on the sphere, we can project it down from the North Pole to give us a map on the plane. Conversely, given any map on the plane, we can project it up onto the sphere.

Notice that these projections do not affect the colouring of the map – if two neighbouring countries are coloured *red* and *green*, then after projection (in either direction) their images will also be *red* and *green*. It follows that we can restate the four-colour problem as a problem about the colouring of maps on a sphere:


Four-colour problem for a sphere

Can every map drawn on the surface of a sphere be coloured with at most four colours in such a way that neighbouring countries are coloured differently?


If we can solve the four-colour problem for maps drawn on a sphere, then we immediately obtain a solution for maps drawn on the plane. Conversely, if we can solve the four-colour problem for maps drawn on the plane, then we immediately obtain a solution for maps drawn on a sphere. So it makes no difference whether we think of our maps as drawn on the plane or on the surface of a sphere and so, for each map, we can choose whether or not to colour the exterior region.

There are some further simplifications which do not materially affect the four-colour problem itself, but which will help to simplify our explanation. It almost goes without saying that every map we consider is in one piece, since we could treat the various pieces as maps in their own right and colour them all separately. Similarly, we can rule out any map that consists of pieces joined together at a single point, since again we could colour the various pieces


separately: in particular, we can ignore countries with a single boundary line. Thus, we can ignore any maps like these:



Finally, we can assume that there are at least three boundary lines at each meeting point. If there were only two boundary lines at a meeting point, then we could simply remove the point without affecting the colouring.



In fact, as we shall see in Chapter 4, when trying to solve the four-colour problem we can always restrict our attention to maps where there are *exactly* three boundary lines at each meeting point, as in the map overleaf. Such maps are very common, and we give them a special name: we call them *cubic maps*. You might like to try to colour this map with four colours. (Can it be coloured with just three?)




TWO EXAMPLES

To end this chapter, here are two practice problems about the colouring of maps on the plane.

Example 1


In the following map, three countries have already been assigned colours. How can we colour the entire map with the four colours *red*, *blue*, *green* and *yellow*?




Notice first that country *A* has a boundary line in common with countries coloured *green* and *yellow*, and so must be coloured either *blue* or *red*. Let us look at each of these possibilities in turn.

If country *A* is coloured *blue*, then *F* must be coloured *red* (since


it shares a boundary line with countries coloured *blue*, *green* and *yellow*). Country *D* must then be coloured *green*, and *E* must be coloured *yellow*. This gives the following colouring.



But it is now impossible to colour country *C*, since it has boundary lines in common with countries of all four colours. So *A* cannot be coloured *blue*, and must be coloured *red*.



Country *F* must now be coloured *blue*, *C* must be coloured *green*, *D* *red*, and *E* *yellow*. We must then colour country *H* *red*, *G* *green*, *B* *yellow*, *I* *green* and *J* *blue*. This completes the colouring.



Example 2

Our second example first appeared as an April Fool's joke! For several years Martin Gardner wrote a highly successful mathematical column in the pages of *Scientific American*, and many of his columns were later collected in book form. In the 1 April 1975 issue of the magazine he decided to play a trick on his readers by presenting 'Six sensational discoveries that somehow or another have escaped public attention'. They included a chess-playing machine (then considered impossible), a thought experiment that disproved Einstein's theory of special relativity, and the discovery of an ancient manuscript establishing Leonardo da Vinci as the first inventor of the valve flush toilet. The column declared:

The most sensational of last year's discoveries in pure mathematics was surely the finding of a counterexample to the notorious four-color-map conjecture . . . Last November, William McGregor, a graph theorist of Wappingers Falls, N.Y., constructed a map of 110 regions that cannot be colored with fewer than five colors.

In the July 1975 issue the April Fool hoaxer owned up, reporting that his column had elicited more than a thousand letters and that

hundreds of readers had sent in colourings of the map with just four colours. McGregor's map appears below. Can you colour it with four colours?

